Material Allocation in MRP with Tardiness Penalties

by

Uday S. Karmarkar

Anderson School of Management
University of California, Los Angeles
Los Angeles, CA

and

Ramakrishnan S. Nambimadom

Weatherhead School of Management
Case Western Reserve University
Cleveland, OH

February 1996

Abstract

In this paper, we address some flaws in the material allocation function of Materials Requirements Planning (MRP). The problem formulation differs from standard MRP logic in certain important ways; start and finish times for orders are forced to be realistic and material allocations are made to minimize the total tardiness penalty associated with late completion. We show that the resulting MRP material allocation problem is NP-hard in the strong sense. A lower bound and a heuristic are developed from a mixed integer linear formulation and its Lagrangean relaxation. The lower bound and the heuristics are closer to the optimum in cases where there is either abundant material or considerable competition for material; in intermediate cases, small perturbations in material allocation can have a significant effect. A group of heuristics based on the MRP approach and its modifications is examined; they are optimal under certain conditions. An improvement method that preserves priorities inherent in any given starting solution is also presented. The Lagrangean heuristic performs better than the MRP based heuristics for a set of 3900 small problems, yielding solutions that are about 5% to 10% over the optimal. The best MRP based heuristic does about as well as the Lagrangean heuristic on a set of 120 larger problems, and is 25% to 40% better than the standard MRP approach, on the data sets tested.

Journal of Global Optimization **9**: 453–482, 1996.

1 Introduction

Materials Requirements Planning (MRP) is probably the most widely used production control software method and has become a *de facto* standard in manufacturing management (Dilworth, 1993). A recent survey of manufacturing managers (Deloitte and Touche, 1990) found that MRP was considered their second most important tool. In its most extensive applications, it carries out or supports a large set of production control functions, including order release (batching and timing), inventory management, material coordination, material allocation, order tracking and data management. Despite the widespread popularity of MRP, it has not been the focus of much research. In particular, the decision problems in MRP are hidden and are solved using simple, and sometimes, arbitrary heuristics. In this paper, we examine the allocation of available material to orders, so as to minimize penalties due to late completion of final item customer orders against due dates. When allocated material is not sufficient to fill a gross requirement generated by an order, it creates a net requirement, which must be satisfied by a production order. This order requires a lead time and generates gross requirements for the next level in the bill of materials (BOM). The lead times involved in production or in obtaining input material may delay the order. Material allocation thus directly affects the completion time of orders at all levels of the BOM, and especially customer orders for end items. It is therefore a central factor in determining on time performance with respect to customer orders.

In many industries today, material constraints can be as important as capacity constraints and hence improved material allocation techniques are as important as capacity management techniques such as sophisticated scheduling methods. For example, over the last few years, there have been frequent component shortages in the computer industry. Computerworld (March 29, 1993) reported that IBM's $ 1 billion backlog for its ThinkPad 700 line of notebook computers can be partly attributed to the shortage of 'literally a nickel part'. The recent trend towards the adoption of lean manufacturing practices has further increased the importance of material allocation. The ability to allocate existing inventories in a judicious fashion decreases the need for large inventories in the first place. A recent survey (Deloitte and Touche, 1990) notes that maintaining delivery performance and controlling inventory costs are two of the key objectives for manufacturers in the 1990s. While these objectives can be viewed as being contradictory to each other, improved methods of allocating existing inventories can allow manufacturers to pursue them simultaneously.

In standard MRP logic, allocation of on hand inventories and scheduled receipts is typically achieved by a fixed decision rule. The most common rule allocates available material to gross requirements in the due date sequence. A more sophisticated rule might first allocate materials to requirements generated by external customer orders for finished items, then allocate to orders for parts or spares, and finally fill orders for stock orders or inventory. In many MRP implementations, the planner has the option of reallocating material by overriding the system's allocations. The system subsequently preserves these allocations. It is clear that any fixed allocation rule of the type described is likely to be non-optimal.

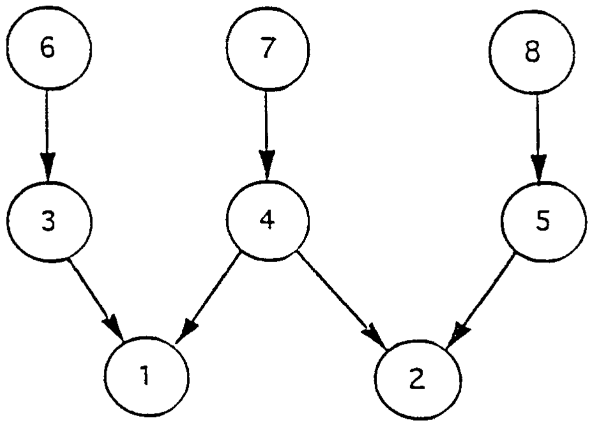

Figure 1: Example 2

It is of some interest to note the different ways in which the standard rule can be sub-optimal. As can be expected, one of the weaknesses of the method used in MRP (the due-date based allocation rule) is that it fails to discriminate between orders on the basis of their 'importance'. In other words, it may be possible to obtain a solution, where an important order would be less tardy than in the solution provided by MRP, albeit, at the cost of a less important order. This may be accomplished by allocating inventory allocated to the latter order by MRP to the former order. This would not be surprising, since standard MRP has no mechanism to recognize order priorities. Interestingly however, it is possible to improve on the allocations made by MRP and obtain solutions where no order is worse off and some are indeed better off. Thus the problems with the material allocation techniques used by MRP go beyond the inability to recognize order priorities; MRP leaves some room for free improvement. For example. consider a bill of material consisting of two items, with item 2 feeding item 1. Item 1 has two end orders. The first is for 15 units and has a due date of 1. The second is for 5 units with a due date of 2. Item 1 and 2 have 10 units each of initial inventory. The lead time for item 1 is 3 units. The due date based rule would allot 10 units of item 1 and 5 units of item 2 to order 1. and 5 units of item 2 to order 2. This would result in order 1 being late by 2 units and order 2 being late by 1 unit. However, if we reverse the sequence in which material is allotted, order 2 would be done in time, while order 1 would be late by the same amount, namely 2 units. Thus total tardiness is reduced by 1 unit. For another example, consider the bill of material shown in figure 1. The first order is for 10 units of item 1 with a due date of 1. The second order is for 10 units of item 2 with a due date of 0.5. Items 1,2,5 and 6 have no initial inventory. Items 3,4,7 and 8 have an initial inventory of 10 units. All lead times are 1 unit. The due date based rule allots 10 units each of items 4 and 8 to order 2 and 10 units of items 3 and 7 to order 2. This results in order 1 being late by 1 unit and order 2 by 1.5 units. Once again a reversal of this sequence enables us to do better. In particular, it results in order 1 being allotted 10 units of items 3 and 4, and hence finishing on time, and order 2 being allotted 10 units of items 7 and 8 and being late by the same 1.5 units as before. While these improvements are easy to spot in this case, in general, we would find a combination of such instances and hence possible improvements may not be easily visible.

In general, it is not hard to devise examples where the performance of such a fixed allocation rule is arbitrarily bad. It is therefore of some interest to examine the possibility of making allocations in some improved way. More fundamentally the role of allocation in MRP is rather hidden, and it is often unclear that there are decision heuristics embedded in the MRP logic. Furthermore, standard MRP implementations do not reveal the consequences of allocation decisions, since the algorithm only makes a backward pass to determine release dates. If release dates are negative, the standard system does not impose non-negativity and work forward to determine the consequences for completion times of individual orders. One aim of this paper is to formalize the decision problem underlying material allocation. The formulation differs from standard MRP by requiring release dates to be non-negative. We also provide conditions under which the method used in MRP provides optimal solutions.

There are not many explicit formulations of the set of problems tackled in MRP. Bitran et al. (1982) consider BOM relationships and integration of planning and release in a discrete period model. Billington et al. (1983) also formulate a discrete period mathematical programming model of planning and batching, including BOM relationships. The goal of these papers is perhaps primarily to explore the integration of capacity planning and order release levels. The aspect of MRP that is not captured very well by such discrete period models is the release process with its lead time offsets. Hackman and Leachman (1989) explicitly address lead times, order offsets, and release timing in developing a general framework for production. Penlesky (1989) examines the problem of maintaining open order due dates in job shops using MRP to plan and schedule their operations. The material allocation problem in relation to order delays and tardiness penalties, is not solved in these papers. Tang(1988) mentions a similar problem of allocating scarce raw materials; the objective function considered there is that of minimizing the maximum tardiness penalty. King (1989) describes IRAM (Intraworks Resource Allocation Model), a PC based optimization tool for allocating material, based on the work by Luss and Smith

(1986). The objective in these papers is to minimize the maximum shortfall. The formulation of this paper is closest to that of Karmarkar (1991), except that here the production lead times are taken as constant and given. However, as in that paper, the formulation involves continuous times, discrete events and discrete orders, and has combinatorial features. It also has a structural relationship to the weighted tardiness sequencing problem.

The formulation analyzed here thus addresses some of the shortcomings in the standard MRP algorithm. However, it does not address some other egregious problems. In particular it is assumed that the production process is approximated by a fixed production (or order) lead time. Furthermore, the issues of batching and safety times are not considered. These are undoubtedly severe simplifications of the real world problem. However, it seems reasonable to extract the material allocation problem as a subject for study for the following reasons. First, little attention has been given till now to formal analyses of MRP techniques. A single study cannot hope to do justice to all the issues involved. Second, the material allocation decision is a significant component of the production plan that is actually implemented, when the MRP procedure is used to release orders. The material allocation decision, in large part, determines the scheduling problem that has to be solved in implementing production. A sensible material allocation plan could be at least as important a determinant of actual performance, as sophisticated scheduling techniques. Finally, the material allocation plan is more visible to a central planner and can be modified more easily than schedules on individual machines. Hence the material allocation decision is an important part of the production control framework. It is worthwhile noting that even the simplified version dealt with here is computationally complex. Relaxing the assumptions made here is not a trivial extension and is a problem left for future study.

We emphasize that the purpose of this study is not to develop optimization methods for application to this class of problems. We would argue that this is not a pragmatic goal; problems encountered in practice are so large that optimal methods would founder in most cases. Rather, our present aim is to develop formal models of MRP to make the imbedded decision problems explicit and to bring them into the research domain so that the function of MRP systems is clearly understood. The eventual purpose is to develop improved practical methods either as enhancements to, or substitutes for, MRP calculations. Given the computational difficulty of the problem, we expect that practical methods will have to be judicious combinations of heuristics. One of the contributions of this paper is to develop a formal model of the material allocation problem. In addition, we also develop methods that perform significantly better than the standard due date based MRP procedure. These procedures include both optimization based methods and modifications of existing procedures. Our continuing efforts in that direction are briefly described in the last section.

In the next section, we set out the notation used in this paper. A very simple version of the problem where the BOM contains only one item, is shown to be NP-complete in the ordinary sense by a reduction from the knapsack problem. Following this, we present a formulation of the problem based on an MRP-like bill of material explosion. A Lagrangean decomposition which yields a lower bound and a heuristic are presented. Four heuristics based on modifications of the standard MRP approach are presented. An improvement procedure that can use any feasible solution as a starting point and attempt to decrease the tardiness of late orders without disturbing the priorities in the initial solution is also developed. Optimal solutions to a restricted class of bills of material are computed using a dynamic programming algorithm. In order to test the quality of the several bounds and heuristics, a set of 3900 small problems with 13 different BOM structures and varying extent of competition for material, is solved. The Lagrangean heuristic out performs the MRP based heuristics, yielding solutions that are 5% to 10% of the optimal. Among the MRP based heuristics, one does better than the rest. These approaches are subsequently tested for a set of 120 larger problems. Here the best MRP based heuristic does as well as the Lagrangean heuristic, yielding solutions that are 25% to 40% cheaper than the standard MRP approach. The standard MRP based heuristic is dominated by all the other heuristics, which is not surprising given that MRP ignores the magnitudes of tardiness penalties.

2 Notation

The notation used in this paper is fairly similar to that of Afentakis et al. (1984). The BOM can be thought of as a directed acyclic graph. Suppose an item j is a direct input to item i, *i. e.* there is an arc in the BOM graph from item j to item i. Then i is said to be an *successor* of j and j is a *predecessor* of item i. Item i is an *ancestor* of item j if, either i is j itself, or there exists a path from j to i in the BOM graph; j is a *descendent* of i, if i is an ancestor of j. Item i is a *raw material* if it has no inputs. A BOM is said to have an *assembly* structure if each item is an input to at most one item. It has a *generalized assembly* structure if, between any two items there exists at most one path; that is, if an item is an input to more than one item, these items have no common ancestors. Clearly an assembly structure is a special case of the generalized assembly structure, where there is only one end item.

In addition to the assumptions stated in the introduction, we shall make the following assumptions to simplify the discussion. As explained below all these assumptions can be relaxed without much difficulty. We assume that end items have no parent items. To accommodate demand for parts and sub-assemblies, we could add dummy nodes in the BOM as parents to such items, with zero lead time. In addition, it is assumed, unless otherwise stated, that the items with no inputs (raw materials) have unlimited initial inventory. This amounts to adding an extra node for each purchased raw material, with the lead time to the parent item being the purchase lead time. We also assume that there are no scheduled receipts at time zero. This again is not a rigid assumption; we could add an extra node for each scheduled receipt with the lead time to the parent being the time at which it would be available. Finally, we shall restrict our discussion to bills of material with a generalized assembly structure. However, as we shall see, most of our models and methods can be extended to problems with general structures.

In addition, let

$\mathcal{I}$ = The set of items in the Bill of Materials graph
$\mathcal{P}_i$ = The set of predecessors of item i
$\mathcal{S}_i$ = The set of successors of item i
$\mathcal{A}_i$ = The set of ancestors of item i
$\mathcal{D}_i$ = The set of descendants of item i
$\mathcal{R}_i$ = The set of raw materials that are descendants of item i
I_j = Initial inventory of item j
L_i = Lead time required to manufacture item i, after all its inputs are available
m_{ji} = The number of units of item j that directly go into one unit of successor i
T_{ji} = The sum of the lead times from item j's parent to ancestor i
M_{ji} = The number of units of item j that go into one unit of ancestor i
$\mathcal{E}$ = The set of end items
N_i = Number of orders of end item i
Q_{ik} = The size of the kth order of end item i
D_{ik} = The due date of the kth order of end item i
w_{ik} = Cost per unit tardiness of the kth order of end item i

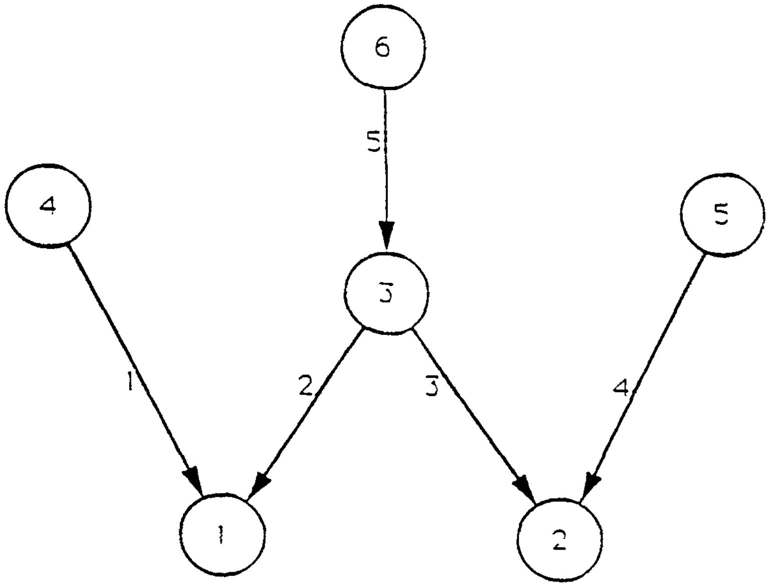

Figure 2: Notation

Numbers alongside arc from j to i represent the number of units of j that go into one unit of i. If j is a direct input to i, j is a *predecessor* of i, and i is a *successor* to j. If either $j = i$ or there exists a path from j to i, j is a *descendent* of i and i is an *ancestor* of j.

Consider Figure 2 to illustrate the notation. The number alongside arc (j, i) in the BOM graph is the quantity of j required by one unit of i (m_{ji}). We shall adopt this convention throughout this paper. Let, $L_1 = 1, L_2 = 2$, and $L_3 = 3$ (there are no lead times associated with items 4, 5 or 6 since, by our convention they are assumed to have unlimited inventory). Then, for example,
$\mathcal{I} = \{1, 2, 3, 4, 5, 6\}$ and $\mathcal{E} = \{1, 2\}$,
$\mathcal{P}_3 = \{6\}$, $\mathcal{S}_3 = \{1, 2\}$, $\mathcal{A}_3 = \{1, 2, 3\}$, and $\mathcal{D}_3 = \{3, 6\}$,
$\mathcal{P}_6 = \phi$, $\mathcal{S}_6 = \{3\}$, $\mathcal{A}_6 = \{1, 2, 3, 6\}$, and $\mathcal{D}_6 = \{6\}$,
$T_{31} = L_1 = 1$, $T_{32} = L_2 = 2$. $M_{31} = m_{31} = 2$, and $M_{32} = m_{32} = 3$,
and,
$T_{63} = L_3 = 3$, $T_{61} = L_3 + L_1 = 4$, $T_{62} = L_3 + L_2 = 5$, $M_{63} = m_{63} = 5$,
$M_{61} = m_{63}m_{31} = 10$, and $M_{62} = m_{63}m_{32} = 15$.

3 Model Complexity

We shall prove that when the BOM contains only one item, the material allocation problem is NP-Hard, by showing that the knapsack problem is a special case. The knapsack problem, which is NP-Hard (Garey and Johnson, (1979)), can be described as follows.
'Given a finite set U, a size $s(u) \in Z^+$ and a value $v(u) \in Z^+$ for each $u \in U$, and positive integers B and K. Is there a subset U' of U, such that $\sum_{u \in U'} s(u) \leq B$ and $\sum_{u \in U'} v(u) \geq K$?'

Now consider the following instance of the material allocation problem. The BOM contains one item, with initial inventory B and a purchase lead time of 1 (as per our notation, a BOM with two items, the second item having infinite initial inventory). For each $u \in U$, we create an order for the item with order size $s(u)$, due date 0, and tardiness cost $v(u)$. Now, if any order u is satisfied from the initial inventory, its tardiness cost is 0; else it is $v(u)$. Therefore we can observe that, if U' is the set of orders that are completely satisfied from initial inventory, it must satisfy the condition, $\sum_{u \in U'} s(u) \leq B$ and the total weighted tardiness cost would be, $\sum_{u \in (U-U')} v(u) = \sum_{u \in U} v(u) - \sum_{u \in U'} v(u)$. Thus we can say that an allocation of initial inventory with total weighted tardiness less than or equal to $\sum_{u \in U} v(u) - K$, exists *iff* the knapsack problem has a solution with value greater than or equal to K. Hence the material allocation problem with a BOM containing only one item is NP-Hard in the ordinary sense.

We can also prove that when the BOM is restricted to be a series, the material allocation problem is NP-Hard in the strong sense. We shall do so by transforming the single machine weighted tardiness problem into a special instance of the material allocation problem with a serial BOM, using a pseudo-polynomial transformation. Since the former problem is NP-Hard in the strong sense (Garey and Johnson, page 237), it implies that the latter is also NP-Hard in the strong sense (Garey and Johnson, page 101).

The single machine weighted tardiness problem ($SMWT$) is,
'Given a set T of jobs, a processing time p_t, a due date d_t and a tardiness cost w_t for each job $t \in T$ and a positive integer K. Is there a single machine schedule with total weighted tardiness K or less ?.'
Let problem MA' be,
'Given a serial BOM, a lead time L_i, initial inventory I_i, and a technology coefficient $m_{i,(i-1)}$ for each item in the BOM, a set of orders with specified order size, due date, and tardiness cost for each order, and a positive integer K. Is there an allocation of inventories, such that total weighted tardiness is less than or equal to K?.'

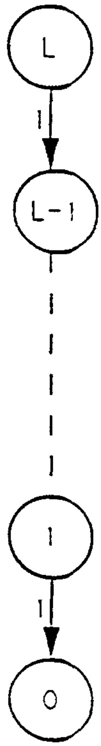

Figure 3: BOM used in complexity proof
BOM obtained by transforming a weighted tardiness scheduling problem with job processing times p_t. $L = \sum p_t$. One unit of an item goes into an unit of it's successor. All lead times are 1. Initial inventory of item 0 is 0; for other items, it is 1. For every job, there is an order of p_t units of item 0, with the same due date and unit tardiness costs.

We shall transform an instance I of problem $SMWT$ into an instance I' of the problem MA'. Consider a serial BOM with $L + 1$ items (figure 3), where $L = \sum_{t \in T} p_t$. Lead times are all 1, initial inventory of item 0 is 0, all other items have an initial inventory of one unit, and one unit of item i goes into an unit of item $i - 1 (i = 1, \ldots, L)$. Corresponding to each job $t \in T$, create an order for item 0, with order size p_t, due date d_t, and tardiness cost w_t. This defines instance I'. This is a pseudo-polynomial transformation, as shown below. In particular, we show that I is a *yes* instance, if and only if, I' is a *yes* instance, and that certain conditions, involving the time required to compute the transformation, the length of instances I and I', and the biggest numbers in I and I', are satisfied. Hence we have shown that MA' is NP-Complete in the strong sense. Since MA' is a special case of the general material allocation problem, the latter is also NP-Complete in the strong sense.

We shall now show that this transformation is a pseudo-polynomial transformation from problem $SMWT$ to MA'. Let Max(I), Max(I') be the biggest numbers in I and I', respectively and let Length(I), Length(I') be the length of I, I', respectively. To show this we need to establish that

1. I is a *yes* instance of $SMWT$ *iff* I' is a *yes* instance of MA',
2. The transformation can be computed in time polynomial in Max(I) and Length(I),
3. There exists a polynomial $q_1 : \forall I, q_1(\text{Length}(I')) \geq \text{Length}(I)$, and
4. Max(I') is bounded by a polynomial in Max(I) and Length(I)

Once we observe that,

$$L = \sum_{t \in T} p_t \leq \text{Max}(I)|T| \leq \text{Max}(I)\text{Length}(I)$$

the last three conditions are obviously satisfied. Clearly the time required to obtain instance I' is polynomial in Length(I) and L, and hence polynomial in Length(I) and Max(I). Condition 3 is easily met since Length(I') > Length(I). And finally,

$$\text{Max}(I') = \max\{L, \text{Max}(I)\} \leq \text{Max}(I)\text{Length}(I)$$

Now we shall prove condition 1. To solve the material allocation problem MA', each order t has to be allocated p_t units. Since each item $i, (i = 1, \ldots, L)$ has exactly one unit of inventory, this means each order has to be assigned a set S_t of p_t items, with the condition that no item is assigned to more than one order. We shall now show that given any solution, we can convert it, without any increase in the objective function. into a sequential solution, that is a solution such that, if $i, j \in S_t$ and $i \leq k \leq j$ then $k \in S_t$. In other words, all orders are allotted items in a sequence. Once we do that. we shall then claim that a sequential solution corresponds to a single machine schedule with the same weighted tardiness cost, thus proving that the transformation is a valid one.

To prove our first claim suppose order 1 is the first one to be finished. Let the last item assigned to it be i. Now assume order 2 is assigned an item $j, (j < i)$. Note that the finish time of order 1 is i and that of order 2 is greater than i. Therefore an interchange of the assignments of items j and i would result in a decrease in the finish time of order 1 and leave the finish time of order 2 unchanged. We can repeat this exchange process, without any increase in the objective function, until we have a sequential solution. Thus we can see that problem MA' has a solution with total weighted tardiness less than K *iff* it has a sequential solution with the same property.

Now if we have a sequential solution to problem MA', the finish time of an order is the sum of its order size p_t, (which is the processing time of the corresponding job) and the order sizes of the orders that were assigned before it. Clearly, if we use the same sequence in the single machine weighted tardiness problem, we shall get the same finish times and the same tardiness costs.

Thus the transformation from instance I of problem $SMWT$ to instance I' of problem MA' is a pseudo-polynomial transformation.

4 Model Formulations

In this section, we present two formulations. The first is a straightforward formulation, based on an MRP-like logic. We then present a second formulation for the problem and show that it is equivalent to the first. The second formulation is presented for the generalized assembly structure; an extension to the general case is discussed. In later sections, we shall use the second formulation for the purpose of development of lower bounds and heuristics.

4.1 A Single Level Peg Based Formulation

This formulation is based on an MRP-like level-by-level approach. An end item order is satisfied from the initial inventory of the end item and internal orders of all its predecessors. An internal order, in turn, is supplied from initial inventory of that item and internal orders of its predecessors. If an order (i, k) is supplied from order (j, l), where j is a predecessor of item i, it is pegged to the order (j, l). Thus this formulation includes binary variables that reflect the pegging status between an order and the internal orders of its predecessors. These pegging variables are used to impose constraints on the finish times of orders. The term order includes both internal and external orders, unless otherwise stated. Order (j, l) refers to the lth order of item j. Let,

$$\begin{aligned}
q'_{jl} &= \text{Quantity allocated from initial inventory of item } j \text{ to order } (j,l) \\
q_{ikjl} &= \text{Quantity allocated from order } (i,k) \text{ to order } (j,l) \ \ \forall i \in \mathcal{P}_j \\
z_{ikjl} &= 1, \text{if } q_{ikjl} > 0, \\
&= 0, \text{otherwise.} \\
Q_{ik} &= \text{Size of order } (i,k) \\
F_{ik} &= \text{Finish time of order } (i,k)
\end{aligned}$$

Note Q_{ik} is given if i is an end item; otherwise it is a decision variable. Let M be a very large number. Then our formulation is,

$$\min_{q,q',z,Q,F} \sum_{i\in\mathcal{E}} \sum_{k=1}^{N_i} w_{ik}[F_{ik} - D_{ik}]^+$$

subject to,

$$\begin{aligned}
\sum_l q'_{jl} &\leq I_j & (1)\\
\sum_{j,l} q_{ikjl} &\leq Q_{ik} \;\; \forall \text{ internal orders } (i,k) & (2)\\
\sum_k q_{ikjl} &= m_{ij}(Q_{jl} - q'_{jl}) \quad \forall i \in \mathcal{P}_j & (3)\\
q_{ikjl} &\leq M z_{ikjl} & (4)\\
F_{jl} &\geq F_{ik} + L_j - M(1 - z_{ikjl}) & (5)\\
q, q', Q, F &\geq 0 & (6)\\
z &= 0 \text{ or } 1 & (7)
\end{aligned}$$

The first two set of constraints ensure that the quantity alloted, from either initial inventory or internal orders, does not exceed the quantity available. The third constraint set ensures that order (j,l) is supplied adequate material from initial inventory and the internal orders of its predecessors. Constraint set 4 ensures that if order (i,k) supplies order (j,l), then the latter is pegged to the former. Finally, the last set of constraints ensure that the finish time of an item's order is at least as much as the finish time of predecessor item orders to which it is pegged, plus the lead time of the item.

Since we need to decide the number of variables, one issue that may arise in using such a formulation, is the number of internal orders that are possible for item i. As there are no batching issues, all material that is available at any time can be grouped together as one internal order. Therefore when the BOM has a generalized assembly structure, the maximum number of internal orders of item i can be set at the number of descendents it has; in the general case it can be set equal to the sum of the number of paths from each descendent. Another approach is to insist that each internal order can supply only one order of its successor, and each order can be supplied from at most one internal order of each of its predecessor items. We can see that these restrictions do not change the finish times of customer orders. The first constraint is clearly no problem, since splitting batches would accomplish this. The second constraint can be satisfied by collecting all the internal orders of a predecessor item that supply an order of a successor item, and combining them into one order. Note that the finish time of the combined order will be the maximum of the finish times of the orders that went into this combined order, and hence the finish time of the successor item's order remains unaffected. Thus we can see that it suffices to have as many internal orders of each item i as the total number of orders of end items that are ancestors of item i. Also note that in this case, internal orders at all levels are associated with an unique end order. The formulation presented in the next section utilises this approach.

4.2 Path Based Formulation

In the previous subsection, we saw that we can, without loss of generality impose a lot for lot policy. This is true, essentially because our model does not consider any setup costs. We also saw that the lot for lot policy implies that each internal order of an item j is associated with an unique end item order. Hence we can refer to an internal order associated with the kth order of end item i as order (j, i, k). If an order is supplied from internal orders of its predecessor items, it is pegged to these orders. Thus the formulation includes binary variables that reflect the pegging status between an order and the internal orders of its predecessors. These pegging variables are used to impose constraints on the finish times of orders.

Let,

$$\begin{aligned}
Q_{jik} &= \text{Size of internal order } (j, i, k) \\
q_{jik} &= \text{Quantity allocated from the initial inventory of item } j \text{ to order } (j, i, k) \\
z_{jik} &= 1, \text{if } Q_{jik} > 0, \\
&= 0, \text{otherwise.} \\
F_{jik} &= \text{Finish time of order } (j, i, k) \\
F_{ik} &= \text{Finish time of order } (i, k)
\end{aligned}$$

The formulation is,

$$\min_{q,z,Q,F} \sum_{i \in \mathcal{E}} \sum_{k=1}^{N_i} w_{ik}[F_{ik} - D_{ik}]^+$$

subject to,

$$\sum_{i \in \mathcal{E}} \sum_{k=1}^{N_i} q_{jik} \le I_j \qquad \forall j \in \mathcal{I} \tag{8}$$

$$Q_{j'ik} = m_{j'j}(Q_{jik} - q_{jik}) \qquad \forall j' \in \mathcal{P}_j \quad \forall j \in \mathcal{D}_i \quad \forall k = 1, \ldots, N_i \quad \forall i \in \mathcal{E} \tag{9}$$

$$Q_{jik} \le (\min(I_j, M_{ji} Q_{ik})) z_{jik} \qquad \forall j \in \mathcal{D}_i \quad \forall k = 1, \ldots, N_i \quad \forall i \in \mathcal{E} \tag{10}$$

$$F_{jik} \geq F_{j'ik} + L_j z_{j'ik} \qquad \forall j' \in \mathcal{P}_j \quad \forall j \in \mathcal{D}_i \quad \forall k = 1, \ldots, N_i \quad \forall i \in \mathcal{E} \tag{11}$$

$$Q_{iik} = Q_{ik} \qquad \forall k = 1, \ldots, N_i \quad \forall i \in \mathcal{E} \tag{12}$$

$$F_{ik} = F_{iik} \qquad \forall k = 1, \ldots, N_i \quad \forall i \in \mathcal{E} \tag{13}$$

$$q_{jik}, Q_{jik} \geq 0 \tag{14}$$

$$F_{jik}, F_{ik} \geq 0 \tag{15}$$

$$z_{jik} = 0 \text{ or } 1 \tag{16}$$

The first constraint set ensures that the quantity allotted from initial inventory is no more than the amount available. The second constraint set ensures that the net requirements of order (j, i, k) are met by the internal orders of item j's predecessors. Constraint sets 10 and 11 enforce the lead time offsetting process, if the net requirement of an item's order is positive. Finally constraints 12 and 13 essentially define order (i, i, k) to be the kth order of end item i.

The formulation can be easily extended to consider the case of bills of material where there may be multiple paths between items. The only difference is that associated with each end item order, the number of internal orders of an item that supplies it is equal to the number of paths. Hence to refer to an internal order we need to include a reference to the path number that the internal order corresponds to. The formulation can then be obtained in an almost identical fashion. (Karmarkar and Nambimadom, 1992).

5 A Lagrangean Relaxation

In this section we present a Lagrangean relaxation of a modified version of our formulation. This relaxation provides a lower bound to the material allocation problem. The multipliers were adjusted using a sub-gradient algorithm (Nemhauser and Wolsey 1988) and the best lower bound was calculated. Heuristic solutions are also obtained as a byproduct of the relaxation. We also briefly discuss a number of other Lagrangean relaxations.

For the purpose of our relaxation it is useful to modify our formulation and replace constraints 11 and 13 by the following.

$$F_{ik} \geq T_{ji} z_{jik} \tag{17}$$

where T_{ji} is the sum of the lead times from item j's parent to item i. To see that this is a valid reformulation, note that by constraint 9, Q_{jik} is greater than 0 implies $Q_{j'ik}$ is greater than 0 for all j' lying on the path between j and i. Hence $z_{jik} = 1$ implies $z_{j'ik} = 1$ for the same j'. Substituting this in constraint set 11 and eliminating the variables F_{jik}, we get constraint 17. Note that we can obtain the reverse transformation by suitably defining the variables F_{jik}.

Relaxing constraint set 8 decomposes the problem by end order. Let λ be a set of Lagrangean multipliers ($\lambda \geq 0$), $z^1(\lambda)$ be the solution to the relaxed problem, and $z^1_{ik}(\lambda)$ be the solution to the sub-problem associated with order (i,k). Then,

$$z^1(\lambda) = \sum_{i \in \mathcal{E}} \sum_{k=1}^{N_i} z^1_{ik}(\lambda) - \sum_{j \in \mathcal{I}} \lambda_j I_j$$

We shall now consider a solution procedure for solving the subproblem associated with order (i,k). For the purpose of the ensuing discussion, ignore items that are not descendants of item i. The subscripts associated with the end item and order number have been dropped, whenever they are unambiguous, for the sake of clarity. Let,

$$\begin{aligned}
\mathcal{D} &= \text{The set of descendants of end item 1} \\
n &= \mathcal{D}| \\
L'_i &= \text{The sum of lead times along the path from item } i \text{ to item 1,} \\
&\quad\ \text{including the lead time for } i \\
&= T_{i1} + L_i \quad \forall i = 1, \ldots, n \\
L'_0 &= 0
\end{aligned}$$

Since we have removed all the items that are not descendants of item 1, $\mathcal{D}$ is now the set of all items. Note that the BOM graph is now a tree. In this section, it is assumed that the items have been renumbered in a non-decreasing sequence of L'_i. Define,

$$\begin{aligned}
Q_j &= \text{The size of the internal order for item } j \quad \forall j \in \mathcal{D} \\
q_j &= \text{The amount of item } j \text{ allocated to order } (1,1) \quad \forall j \in \mathcal{D} \\
z_j &= 1, \text{ if } q_j > 0 \\
&= 0, \text{ otherwise} \qquad \forall j \in \mathcal{D} \\
F &= \text{Finish time of order } (1,1)
\end{aligned}$$

Then the sub-problem associated with order (1,1), $P_{11}(\lambda)$, is,

$$z^1_{11}(\lambda) = \min_{q,z,F} w[F - D]^+ + \sum_{j \in \mathcal{D}} \lambda_j q_j$$

subject to,

$$Q_{j'} = m_{j'j}(Q_j - q_j) \quad \forall j' \in \mathcal{P}_j \ \forall j \in \mathcal{D} \tag{18}$$

$$
\begin{aligned}
Q_j &\leq (\min\{I_j, M_{j1}Q\})z_j \quad \forall j \in \mathcal{D} && (19)\\
F &\geq T_{j1}z_j \quad \forall j \in \mathcal{D} && (20)\\
Q_1 &= Q && (21)\\
Q_j, F &\geq 0 && (22)\\
z_j &= 0 \text{ or } 1 && (23)
\end{aligned}
$$

The above problem is a mixed integer program. The Lagrangean multipliers associated with an item can be interpreted as material costs. The objective is to minimize the sum of tardiness and material costs, while ensuring that the material needed by the customer order is satisfied.

The integral variables in the above problem are a result of the tardiness cost term. If we set the finish time variable F at a particular value, the tardiness cost is fixed. In addition. material can be allotted from only those items which guarantee that the finish time of the customer order is no more than the value of F (constraint 20). The minimum total material cost is now given by a linear program. We can also see that the $n+1$ possible values of F are given by $L'_0, \ldots, L'_n$. The tardiness costs is non-decreasing and the total material cost is non-increasing as we increase F from L'_0 to L'_n. Thus we can solve the above problem by solving $n+1$ linear programs and picking the minimum total cost solution. Solving $n+1$ linear programs using a standard LP code can be time consuming. Note that the only difference in the linear programs resulting from fixing F at L'_j and fixing it at L'_{j+1} is that in the latter, Q_j can be greater than 0 for the predecessors of item $j+1$. Appendix B provides an alternative algorithm which exploits this feature. Another alternative is to use the optimal basis for the former linear program as the starting basis of the latter. In this paper the first approach has been taken. The resultant algorithm has a complexity of $O(n^3)$ for solving the sub-problem associated with an end order.

Karmarkar and Nambimadom (1992) discuss a number of other Lagrangean relaxation based procedures to generate lower bounds and heuristics. However their numerical experiments show that the methods reported in this paper are superior. They also provide a dynamic programming algorithm that calculates the optimal solution for small problems.

6 Heuristics

We developed a number of heuristics for the material allocation problem. They can be divided primarily into two groups. The first group of heuristics is based on the Lagrangean relaxation examined in Section 5 or those suggested there. The second group of heuristics is based on modifications of the standard MRP based approach.

Finally, we also obtained an improvement heuristic which starts with any initial solution and attempts to make it better. We shall first discuss the various heuristics and then describe the heuristics whose computational performance is reported in this paper.

6.1 Lagrangean heuristics

The Lagrangean heuristic tested in this paper uses the solution to the relaxed problem with the multipliers that yielded the best lower bound as a starting point. Customer orders are sequenced by the earliest finish time in this solution and material is assigned to orders as per the sequence. Karmarkar and Nambimadom (1995) examines a number of other approaches to generate heuristics. However as in the case of lower bounds, the methods used in this paper prove to be superior.

6.2 MRP based heuristics

The simplest MRP based heuristic allocates initial inventory according to the earliest due date rule. Recall that the due date here refers to the due date of the internal order (net requirement), obtained by the lead time off-setting process. Thus material is allotted in the sequence in which net requirements are created. Internal orders are generated for predecessor items to meet the net requirements. The due dates for the internal orders are obtained by offsetting the original order's due date by the lead time needed to assemble the predecessor items. This process could end up with negative start times. In that case, start times are pushed to 0 and a forward pass is made to calculate the actual finish times and the resultant tardiness. If all the start times are nonnegative at the end of the backward pass, all the orders can be done in time and thus the optimal solution, which is given by this heuristic, has an objective function value of 0. In fact, as shown by the following proposition, the optimal objective function value can be 0 only under this scenario. This result establishes the optimality of the MRP heuristic if there exists a solution in which all orders can be completed in time.

Proposition 1 *If in the solution given by the MRP heuristic, at least one order is delayed, then there exists no solution in which all orders can be completed in time.*

Proof
Consider an order which is tardy in the solution given by the MRP heuristic. Suppose it is tardy because it is allotted material from the initial inventory of an item j.

To finish this order by its due date, one has to allot the material allotted from an ancestor of item j to another order, to this order. This may not be possible because of insufficient initial inventory of the ancestors of item j. In that case our claim is true. If not, the order which was originally allotted this material has to be allotted material from item j's inventory. But since this latter order had priority over the tardy order in the solution given by the due date based MRP heuristic, it must have an even tighter due date and hence will now be tardy. Thus it is not possible to complete the tardy order in time without making some other order late. □

We consider two modifications to this heuristic. In both, material is allotted to orders on the basis of the earliest due date rule, if no order is already tardy. If not, all orders with a negative due date are given priority over the others. Orders with negative due dates are sequenced according to their tardiness penalties (w_{ik}) in the first modified heuristic and according to the ratio of the tardiness penalty to the order size in the second modified heuristic (order size refers to the size of the internal order). The orders with positive due dates are sequenced by the earliest due date rule. Both these modifications of the simple MRP heuristic share its property of finding the optimal solution, if it is possible to do all jobs in time.

6.3 An improvement heuristic

An improvement heuristic which can be used in conjunction with any other heuristic can be obtained as follows. In both the examples discussed in the introduction, the improved performance was obtained by recognizing the inevitability of one of the jobs being delayed. Based on this observation, material allotment was adjusted so as to make an effort to meet the due date of the other order, without further delaying the tardy order. Thus given any material assignment, we could treat the finish times in this assignment as the due date and apply the simple due date based MRP heuristic, with ties broken randomly. One could repeat this process till there is no improvement in the objective function. The following proposition guarantees the convergence of this heuristic.

Proposition 2 *In each iteration, the finish time of an order is no more than its finish time in the previous iteration.*

Proof
This follows from the following observations. First, note that the finish time in the previous iteration is treated as the due date for the current iteration. Hence there exists a solution where all orders can be completed by their current due date

(the solution obtained in the previous iteration is one possibility). Second, as a consequence of Proposition 1, the MRP heuristic will also provide a solution where all orders can be completed by their due date. Thus the finish time of an order is no more than its finish time in the previous iteration. □

Proposition 2 also proves that the improvement heuristic does not change improve any order's performance at the cost of another order. This can be an important consideration if the starting solution reflected relative priorities among orders that have been set by another algorithm or by managerial intervention. Another interesting feature of the improvement heuristic is that it does not require explicit specification of tardiness penalties, since it preserves the relative priorities implicit in the starting solution. Thus the only information needed is the data in the Bill of Materials.

The heuristics whose performance is reported in this paper are LH, the first Lagrangean heuristic. augmented by the improvement procedure; $MRP1$, the standard due date based MRP heuristic; $MRP2$, which is $MRP1$ followed by the improvement heuristic; $MRP3$. the first of the modified MRP heuristics (where orders with negative due dates are sequenced in decreasing order of their tardiness penalties, w_{ik}), with the improvement heuristic; and finally, $MRP4$, which is same as $MRP3$, except that orders with negative due dates are sequenced by the ratio of their tardiness penalties to their order sizes.

7 Computational Results

We first tested our procedures for 13 small Bills of Material, and later on 4 larger ones. For each of the 13 small bills of material, 300 sample problems were solved. 30 sample problems were solved for the larger bills of material. Table 1 summarizes some of the characteristics of the BOM in the data set. Further details are provided in Karmarkar and Nambimadom (1992). For the first 13 BOM. the optimal solution was obtained using the DP algorithm. We evaluated the lower bound obtained by the Lagrangean relaxation (LB) and the five heuristics described above ($LH, MRP1, MRP2, MRP3$. and $MRP4$).

Table 1: Characteristics of the BOM structures used in numerical tests

BOM	Items	End items	Levels	Width		Paths to end items	
				Mean	Max	Mean	Max
1	3	1	3	1.00	1	1.00	1
2	3	1	2	1.50	2	1.00	1
3	4	1	2	2.00	3	1.00	1
4	4	1	4	1.00	1	1.00	1
5	5	1	3	1.67	2	1.00	1
6	3	2	2	1.50	2	1.33	2
7	4	2	2	2.00	2	1.50	2
8	5	2	2	2.50	3	1.20	2
9	5	2	2	2.50	3	1.60	2
10	5	2	3	1.67	2	1.40	2
11	6	2	3	2.00	3	1.33	2
12	6	3	2	3.00	3	1.50	2
13	7	3	2	3.50	4	1.29	2
14	31	1	5	6.20	16	1.00	1
15	35	3	4	8.75	16	1.23	2
16	30	1	9	3.33	5	1.00	1
17	35	2	6	5.83	8	1.49	2

BOM : Bill of material number; Items : Number of items; End items : Number of end items; Levels : Number of levels in the bill of material; Width : Number of items on the same level: Paths to end items : Number of paths leading from an item to end items

We will start our description of the sample problems by first providing an overview of the data generation procedure and then the details of each step. The initial inventory data was generated randomly, using a procedure which attempts to ensure that material is uniformly spread, as we go from the roots to the raw materials. After the initial inventory data is generated, the customer order data is generated. The procedure that generates customer order sizes utilizes the initial inventory data and a parameter termed the Material Competition Factor (MCF). By varying the MCF value, we can adjust the degree of competition for material. In particular, a MCF value of 1 indicates that there is a rough balance between total inventories and total customer orders; a higher value indicates scarcity of inventory and a lower value, abundance. Due dates for customer orders do not exceed the total lead time for purchasing and production of the end items.

Each end item's initial inventory was uniformly distributed between 25 and 125. For each path leading from an item j to an end item i, a random number between 25 and 125 was generated. This number was then multiplied by the number of units of item j, that went into one unit of end item i along that path. The initial inventory of item j, was taken to be the sum of these products, over all paths leading from j to end items. To generate the order sizes for the customer orders of an end item, first the longest path leading to that end item was identified. Then for each item on this path, we calculated the number of end item units that could be produced using that item's initial inventory, provided there was sufficient inventory of items not on this path. In other words, we divided the initial inventory by the number of units of that item needed to produce one end item. These were then added up over items on the longest path (excluding the leaf item, which by our convention has unlimited inventory). Let this sum be A. Thus A is the number of units of the end item that could be produced, if there were abundant supply of items not on the longest path. Having calculated this, we let the order size be an uniform random variable between 1 and $(2A \times MCF/N_i)$, where MCF is a parameter which controls the competition for material. For generating the due dates, the maximum of the sum of lead times along paths leading from raw materials to the end item was calculated. The due dates of customer orders were uniformly distributed between 0 and this quantity. The tardiness costs were uniformly distributed between 0 and 1.

For each bill of material, we tested the lower bounds and the heuristics for three values of the Material Competition Factor (0.5, 1.0 and 2.0). The number of end orders for each end item was set at 10 for BOMs with a single end item, 5 for those with two end items and 4 for three end item cases. Thus in the one or two end item cases there were 10 end orders, and 12 in the three end item BOMs. For each BOM, MCF combination, 100 problem instances were solved. After generating an instance the optimum solution was first calculated. If it turned out to be 0, the instance was discarded, and a new instance was generated. Otherwise the lower bound (LH) and heuristic solutions ($LH, MRP1, MRP2, MRP3$, and $MRP4$), were calculated as fractions of the optimal solution. Then the mean and standard deviation of these values over the 100 instances were calculated. The procedures were coded in Pascal (Turbo Pascal Version 7) and run on a Gateway 2000 486/66 MHz PC.

These results are tabulated in Table 2 with the sample standard deviation in brackets. The performance of the different methods did not seem to exhibit a dependence on the BOM structure used. So we aggregated the results over all Bills of Material to obtain the first set of numbers depicted in Table 4. A close examination of the results indicated that the lower bound and the heuristics often 'performed poorly' when the optimal objective function value was small. This turned out to be due to the fact that we judge the solutions obtained by the lower bound and the heuristics on the basis of the ratio of their objective function value to the optimal objective function value. When the latter is small, this ratio appears to be large, even if the actual difference between, say, the lower bound and optimal is small. So to judge the performance of the different procedures fairly, we separated out the sample problems on the basis of the optimal objective function value. If this value was more than 1, the problem was deemed to be 'easy', otherwise it was deemed to be 'difficult'. The

Table 2: Results for BOM 1-13, as a Fraction of the Optimal Solution
(see Table 3 for explanation of terms)

BOM	MCF	LB	LH	MRP1	MRP2	MRP3	MRP4
1	0.5	0.69 (0.32)	1.20 (0.96)	3.57 (13.19)	3.47 (13.19)	3.39 (13.20)	3.37 (13.20)
	1.0	0.63 (0.34)	5.79 (37.27)	8.49 (38.93)	8.39 (38.93)	8.25 (38.95)	7.87 (38.29)
	2.0	0.81 (0.17)	1.09 (0.30)	2.22 (1.47)	2.06 (1.39)	1.95 (1.34)	1.91 (1.35)
2	0.5	0.70 (0.26)	1.18 (0.87)	3.85 (12.70)	3.71 (12.71)	3.68 (12.71)	3.66 (12.71)
	1.0	0.74 (0.27)	1.39 (2.78)	2.72 (3.58)	2.61 (3.60)	2.52 (3.62)	2.48 (3.63)
	2.0	0.87 (0.11)	1.07 (0.21)	1.97 (1.06)	1.85 (1.00)	1.67 (0.83)	1.63 (0.79)
3	0.5	0.69 (0.28)	1.28 (1.58)	2.61 (2.68)	2.47 (2.63)	2.40 (2.63)	2.38 (2.63)
	1.0	0.68 (0.31)	1.52 (2.41)	2.78 (3.54)	2.59 (3.17)	2.52 (3.18)	2.50 (3.19)
	2.0	0.82 (0.18)	1.03 (0.06)	2.17 (1.39)	2.02 (1.21)	1.88 (1.10)	1.84 (1.11)
4	0.5	0.65 (0.37)	1.29 (1.10)	1.53 (1.20)	1.51 (1.20)	1.50 (1.20)	1.50 (1.20)
	1.0	0.61 (0.38)	1.38 (0.98)	2.30 (2.21)	2.17 (2.05)	2.03 (2.02)	2.02 (2.02)
	2.0	0.81 (0.22)	1.12 (0.52)	2.08 (1.16)	2.00 (1.15)	1.69 (1.07)	1.59 (0.83)
5	0.5	0.69 (0.31)	1.80 (5.55)	3.92 (10.03)	3.82 (10.04)	3.70 (10.05)	3.72 (10.05)
	1.0	0.58 (0.35)	4.17 (21.25)	5.02 (21.16)	4.91 (21.17)	4.76 (21.19)	4.75 (21.19)
	2.0	0.84 (0.16)	1.18 (1.07)	1.93 (1.22)	1.88 (1.21)	1.74 (1.18)	1.66 (1.18)
6	0.5	0.69 (0.29)	1.07 (0.34)	2.26 (2.78)	2.11 (2.68)	2.11 (2.68)	2.11 (2.68)
	1.0	0.71 (0.29)	1.12 (0.64)	1.65 (1.02)	1.56 (0.99)	1.50 (0.98)	1.50 (0.98)
	2.0	0.82 (0.17)	1.17 (1.22)	2.01 (1.52)	1.82 (1.48)	1.65 (1.41)	1.64 (1.42)
7	0.5	0.70 (0.34)	1.29 (1.13)	1.66 (1.37)	1.59 (1.33)	1.56 (1.30)	1.56 (1.30)
	1.0	0.69 (0.30)	1.43 (1.70)	2.28 (2.32)	2.04 (1.87)	2.04 (1.86)	2.04 (1.87)
	2.0	0.82 (0.16)	1.13 (0.33)	2.07 (0.96)	1.89 (0.85)	1.83 (0.84)	1.83 (0.85)
8	0.5	0.68 (0.37)	1.13 (0.49)	1.58 (1.85)	1.50 (1.82)	1.48 (1.82)	1.49 (1.82)
	1.0	0.73 (0.29)	1.14 (0.40)	1.70 (0.93)	1.57 (0.68)	1.55 (0.67)	1.53 (0.66)
	2.0	0.89 (0.09)	1.05 (0.15)	1.84 (0.77)	1.68 (0.72)	1.67 (0.73)	1.66 (0.73)
9	0.5	0.67 (0.32)	1.11 (0.55)	2.74 (4.65)	2.66 (4.65)	2.64 (4.66)	2.63 (4.66)
	1.0	0.63 (0.28)	2.10 (9.38)	2.95 (9.45)	2.83 (9.47)	2.79 (9.47)	2.76 (9.45)
	2.0	0.84 (0.14)	1.05 (0.09)	1.71 (0.62)	1.58 (0.57)	1.47 (0.60)	1.44 (0.56)
10	0.5	0.74 (0.28)	1.10 (0.35)	1.96 (3.96)	1.92 (3.97)	1.89 (3.97)	1.88 (3.97)
	1.0	0.65 (0.30)	1.36 (2.22)	2.23 (3.68)	2.01 (3.16)	1.90 (3.15)	1.90 (3.15)
	2.0	0.83 (0.17)	1.06 (0.16)	1.84 (0.86)	1.68 (0.69)	1.49 (0.66)	1.46 (0.65)
11	0.5	0.70 (0.35)	1.13 (0.69)	3.84 (20.11)	3.78 (20.12)	3.60 (20.12)	3.61 (20.12)
	1.0	0.61 (0.32)	2.02 (7.07)	2.78 (7.40)	2.59 (7.30)	2.51 (7.30)	2.51 (7.30)
	2.0	0.85 (0.14)	1.05 (0.11)	1.89 (1.05)	1.69 (0.91)	1.44 (0.52)	1.41 (0.51)
12	0.5	0.73 (0.29)	1.06 (0.29)	1.62 (1.62)	1.38 (1.24)	1.37 (1.24)	1.36 (1.24)
	1.0	0.67 (0.30)	1.24 (0.98)	1.86 (1.39)	1.74 (1.19)	1.63 (1.13)	1.65 (1.14)
	2.0	0.83 (0.14)	1.11 (0.43)	1.76 (0.77)	1.58 (0.62)	1.49 (0.59)	1.46 (0.58)
13	0.5	0.69 (0.31)	1.18 (1.22)	1.82 (1.98)	1.67 (1.96)	1.65 (1.95)	1.66 (1.95)
	1.0	0.65 (0.27)	1.17 (0.54)	1.85 (1.41)	1.67 (1.23)	1.59 (1.18)	1.58 (1.09)
	2.0	0.77 (0.17)	1.06 (0.10)	1.53 (0.52)	1.38 (0.47)	1.34 (0.46)	1.32 (0.47)

next two set of numbers in Table 4 show the performance of the different procedures when we aggregate separately over 'easy' and 'difficult' problems. Table 5 shows the average absolute difference between the objective function values obtained by the lower bound and the various heuristics and the optimal objective function value.

A number of interesting results are apparent. First, though the performance of the lower bound and the heuristics appears to be relatively poor for the 'easy' problems, the actual magnitude of the difference from the optimal (as indicated in Table 5) is smaller for these problems. In other words, the lower bounds and the heuristic do not really do badly on the 'easy' set of sample problems; they just appear to do so, because we look at the ratios relative to the optimal. Our comments below about the lower bounds and the heuristic solutions will be confined to their performance on the 'difficult' set of problems. Second, we can see that the Lagrangean heuristic provides solutions that are within about 5% of the optimal when the Material Competition Factor is 0.5 or 2.0. However when MCF is equal to 1, the Lagrangean heuristic is within 10 % of the optimal. We found that the Lagrangean heuristic provided the optimal solution in about half of the problems in the 'difficult' data set. The various MRP based heuristics are clearly inferior to the Lagrangean heuristic. The standard MRP heuristic yields solutions that are, on average, 155% to 178% of the optimal, as MCF increases from 1 to 4. The improved MRP heuristics do better; they are usually within about 140 % to 150 % of the optimal. The lower bound is, on average, between 80 % to 85 % of the optimal. As in the case of the Lagrangean heuristic, the lower bound does better for MCF values of 0.5 and 2.0 than for the intermediate value of 1.

The performance of the four heuristics and the lower bound was also examined for four larger bills of material with more customer orders. However, the optimal solution could not be calculated due to the size of these problems. Hence we only present the ratios of the lower bound and the heuristics with respect to the standard MRP heuristic ($MRP1$). The number of orders was set at 50 for single end item cases, 25 per end item (total of 50) for two end item cases and 20 per end item (total of 60) for three end item cases. 10 sample problems were generated for each BOM, MCF combination. The mean and standard deviation of these ratios over the 10 problems for each set were calculated. Table 3 shows the results.

Table 3: Results for BOM 14-17, as a Fraction of the MRP1 Heuristic

BOM	MCF	LB	LH	MRP2	MRP3	MRP4
14	0.5	0.12 (0.10)	0.75 (0.22)	0.97 (0.02)	0.82 (0.15)	0.81 (0.16)
	1.0	0.39 (0.08)	0.59 (0.13)	0.98 (0.02)	0.75 (0.07)	0.64 (0.08)
	2.0	0.58 (0.07)	0.73 (0.09)	0.99 (0.02)	0.85 (0.07)	0.76 (0.06)
15	0.5	0.15 (0.06)	0.76 (0.12)	0.95 (0.05)	0.78 (0.14)	0.78 (0.14)
	1.0	0.41 (0.05)	0.65 (0.07)	0.95 (0.03)	0.78 (0.11)	0.67 (0.09)
	2.0	0.59 (0.05)	0.75 (0.06)	0.97 (0.02)	0.89 (0.04)	0.76 (0.07)
16	0.5	0.15 (0.06)	0.72 (0.16)	0.96 (0.04)	0.74 (0.13)	0.73 (0.12)
	1.0	0.42 (0.07)	0.66 (0.07)	0.98 (0.02)	0.70 (0.10)	0.61 (0.10)
	2.0	0.60 (0.07)	0.70 (0.10)	0.99 (0.01)	0.85 (0.09)	0.70 (0.07)
17	0.5	0.13 (0.05)	0.66 (0.25)	0.90 (0.06)	0.66 (0.13)	0.65 (0.14)
	1.0	0.41 (0.05)	0.60 (0.06)	0.96 (0.04)	0.67 (0.10)	0.60 (0.07)
	2.0	0.59 (0.06)	0.69 (0.04)	0.99 (0.01)	0.82 (0.09)	0.71 (0.07)

Table shows mean ratios over 10 instances (100 instances for tables 2, 4 and 5); brackets contain sample standard deviations. MCF : Material Competition Factor; LB : lower bound obtained by the Lagrangean Relaxation; LH : heuristic obtained from the Lagrangean relaxation with improvement heuristic; MRP1 : due date based MRP heuristic; MRP2 : MRP1 with improvement heuristic; MRP3, MRP4 : Modified MRP heuristics with improvement heuristic

Table 4: Summary of Results for BOM 1 to 13

Type of Problems	MCF	LB	LH	MRP1	MRP2	MRP3	MRP4
ALL	0.5	0.69	1.22	2.54	2.43	2.38	2.38
	1.0	0.66	1.99	2.97	2.82	2.74	2.70
	2.0	0.83	1.09	1.92	1.78	1.64	1.60
EASY	0.5	0.60	1.33	3.21	3.09	3.05	3.06
	1.0	0.47	3.24	4.65	4.46	4.41	4.34
	2.0	0.61	1.50	3.49	3.13	2.98	2.92
DIFFICULT	0.5	0.82	1.05	1.55	1.47	1.40	1.39
	1.0	0.80	1.09	1.76	1.64	1.53	1.51
	2.0	0.85	1.05	1.78	1.65	1.52	1.48

Table shows the performance of the various procedures over 1300 sample problems for each MCF value. The first set of numbers (ALL) indicate performance over the entire set of 1300 sample problems. The last two set of numbers (EASY and DIFFICULT) show performance when data set is divided up into two sets based on the optimal objective function value. The EASY data set contains problems for which the optimal objective function value is less than 1.0, the DIFFICULT data set contains the rest. See Table 3 for description of the procedures.

Table 5: Summary of Results for BOM 1 to 13 - Absolute Differences with Optimal

Type of Problems	MCF	LB	LH	MRP1	MRP2	MRP3	MRP4
	0.5	0.28	0.10	0.93	0.80	0.68	0.65
ALL	1.0	0.50	0.22	1.93	1.62	1.30	1.24
	2.0	0.79	0.27	4.15	3.46	2.62	2.29
	0.5	0.13	0.06	0.37	0.32	0.30	0.30
EASY	1.0	0.21	0.17	0.63	0.55	0.52	0.51
	2.0	0.20	0.15	1.34	1.10	1.01	0.97
	0.5	0.50	0.16	1.73	1.49	1.22	1.15
DIFFICULT	1.0	0.71	0.26	2.87	2.39	1.87	1.77
	2.0	0.84	0.28	4.41	3.67	2.76	2.41

Table shows the averages of the absolute deviation from the optimal objective function value (as opposed to the ratios shown in Tables 2, 3 and 4). See Tables 3 and 4 for descriptions of the heuristics and data sets, respectively.

For these problems, $MRP4$ did about as well as the Lagrangean heuristic. Among MRP heuristics the sequence remained the same as before. The best heuristics (LH and $MRP4$) were about 25 % to 40 % better than the standard MRP heuristic. The best performance was obtained for MCF values of 1.0. The lower bound was about 15%, 40% and 60% of the solution given by the MRP heuristic for MCF values of 0.5, 1.0, and 2.0, respectively. The gap between the lower bound and the heuristics is large, particularly for low values of MCF. Judging from our experience with the smaller problems, the sub-optimality of the MRP heuristic is probably the bigger contributory factor.

Running times are reported in Karmarkar and Nambimadom. The improvement procedure typically does 3 passes, and hence the running time for $MRP2$ is three times that of $MRP1$. The modified MRP heuristics by themselves take about 50% more time than the standard MRP approach. However the program developed was not optimized with respect to running time; exploitation of more sophisticated data structures is likely to narrow the gap between the standard MRP approach and the other procedures discussed here, to a considerable extent. The Lagrangean relaxation took about 2000 times the running time of the MRP procedure. However our experience suggests that storing some intermediate calculations can decrease the running times by a factor of 50. Furthermore, the number of iterations in the relaxation can be decreased considerably without seriously affecting the quality of the Lagrangean heuristic (though the lower bound may be significantly looser). To summarize, obtaining the optimal solution appears to be unrealistic for large problems; however the methods suggested in this paper, particularly $MRP4$ and the Lagrangean heuristic with the modifications suggested above, perform significantly better than the standard MRP approach at the cost of a moderate increase in running time.

8 Conclusions and Future Research

In this paper, we have studied the material allocation problem in MRP systems. This problem was shown to be NP-Complete. A mixed integer programming formulation was presented and a Lagrangean relaxation was developed. The relaxation yields a lower bound and a heuristic. Four heuristics based on the due date based MRP approach and its modifications were also developed. It was shown that the heuristics are guaranteed to produce the optimal solutions in certain important cases. In addition, a procedure designed to improve any existing solution, while preserving the priorities inherent in the solution was examined. This procedure does not require explicit specification of tardiness penalties. An optimal algorithm was developed for a class of problems. A set of 3900 small problems involving 13 bills of material was used in an initial test. In this test, the ratio of the lower bound, the upper bound and the four MRP heuristics to the optimal solution was calculated. The Lagrangean heuristic dominated the MRP based heuristics; its solutions were within 5% to 10% of the optimal solution. The heuristic and the lower bound performed better, relative to the optimal, when material was either scarce or abundant than the case with intermediate material availability. Among the 4 MRP heuristics one generally performed better than the others. A test over a set of larger problems involving 4 bills of material was also carried out. In this case the Lagrangean heuristic and the best MRP based heuristic gave similar results, with costs about 25% to 40% less than the standard MRP heuristic. One particularly interesting conclusion is that the modified MRP procedures yield significant improvements at a moderate computational cost.

One avenue of future research would be to develop improved computational procedures for this problem. Another area which needs to be examined, deals with the nature of the problems examined. In this paper, inventory was generally spread out equally through the system. Thus the issues involved were generally those of discriminating between orders based on due dates, tardiness penalties and order size. The MRP approach addresses the first issue to some extent, but not the later two aspects. It would be of interest to examine situations where the amount of inventory present along various paths varied significantly. It would also be of considerable interest to provide the MRP system with some ability to consider scheduling issues. Work is now in progress on developing a MRP system with variable lead times. In addition, a project underway examines the integration of material planning and detailed scheduling.

Acknowledgment

The authors wish to thank Professor Gregory Dobson for suggestions on the complexity proof in Karmarkar and Nambimadom (1992).

Appendix A

Details of the Complexity Proof

Appendix B

An Algorithm for solving the Lagrangean Relaxation (Section 5)

Suppose the finish time, F is fixed at L'_j. Then let $C_{j'}(q)$ represent the cost of obtaining the qth unit of item j' in the cheapest manner possible, using only the inventories of those items that will allow the order to be finished by time L'_j. Suppose j' is an allowed item. If it is a raw material, there is no material cost associated with it, and hence $C_{j'}(q)$ is 0 for all q. If j' is greater than j, then its predecessors cannot be used; else the finish time of the order will be at least $L'_{j'}$, which is greater than L'_j by our numbering convention. Hence $C_{j'}(q)$ will be $\lambda_{j'}$, for $q \leq I_{j'}$, and infinity for larger values. If $j' \leq j$, it can be assembled from its predecessors. Let $C'_{j'}(q)$ be the cost of obtaining the qth unit of item j', in the cheapest manner possible by assembling its predecessors (in other words, we do not consider the possibility of using the initial inventory of item j'). This cost would be the sum of the cost of obtaining each of the predecessors of j' in the required quantities. Therefore,

$$C'_{j'}(q) = \sum_{j'' \in \mathcal{P}_{j'}} \sum_{q'=m_{j''j'}(q-1)+1}^{m_{j''j'}q} (C_{j''}(q')$$

Now to obtain, $C_{j'}(q)$, we note that the inventory of item j' would be used if it has not been completely used and it is cheaper than assembling the predecessors. Therefore, if q' is such that,

$$C'_{j'}(q'-1) < \lambda_{j'} \leq C'_{j'}(q')$$

we can set,

$$\begin{aligned} C_{j'}(q) &= C'_{j'}(q), \quad q < q' \\ &= \lambda_{j'}, \qquad q' \leq q < q' + I_j - 1 \\ &= C'_{j'}(q - I_j), \quad q \geq q' + I_j \end{aligned}$$

Note that $C_{j'}(q)$ is piecewise constant. The number of discontinuities is equal to the number of descendants of item j' that are allowed. This observation can be used to calculate this function in a time proportional to the number of descendants of item j'. The material cost of allotting material to the customer order is given by $\sum_{q=1}^{Q} C_1(q)$. To calculate the cost function associated with the end item we need to recalculate the cost functions of all items lying on the path between j and the end item. Finally, F has to be fixed at $n+1$ different values. Thus, the overall complexity of this algorithm is $O(n^3)$. Note that once we solve the problem with the finish time F fixed at L'_{j-1}, we can use the same cost functions for all items that do not lie on the path between j and the end item, to solve the problem with F fixed at L'_j.

References

[1] Afentakis, P., B. Gavish and U. Karmarkar, "Computationally Efficient Optimal Solutions to the Lot Sizing Problem in Multistage Assembly Systems," *Management Science*, 30. 2 (1984), 222-239.

[2] Billington, P. J., J. O. McClain and L. J. Thomas, "Mathematical Programming Approaches to Capacity Constrained MRP Systems : Review, Formulation and Problem Reduction," *Management Science*, 29, 10 (1983), 1126-1141.

[3] Bitran, G. R., E. A. Haas and A. C. Hax, "Hierarchical Production Planning : A Two-Stage System," *Operations Research*, 30, 2 (1982), 232-251.

[4] Dilworth, J. B., *Operations Management : Design, Planning and Control for Manufacturing and Services*, McGraw-Hill, (1993).

[5] Garey, M. R. and D. S. Johnson, *Computers and Intractability : A Guide to NP-Completeness*, Freeman, (1979).

[6] Hackman, S. T. and R. C. Leachman, "A General Framework for Modelling Production," *Management Science*, 35, 4 (1989), 478-495.

[7] Held, M. and R. M. Karp, "A Dynamic Programming Approach to Sequencing Problems," *Journal of the Society for Industrial and Applied Mathematics*, 10, 1 (1962), 196-210.

[8] Karmarkar, U. S., "Hierarchical Decompositions of Closed Loop MRP Systems," Working Paper, William E. Simon School of Business Administration, (1991).

[9] Karmarkar, U. S. and R. S. Nambimadom, "Material Allocation in MRP with Tardiness Penalties," Working Paper QM 91-06, William E. Simon School of Business Administration, (1991).

[10] King, J. H., "Allocation of Scarce Resources in Manufacturing Facilities," *AT&T Technical Journal*, 68, 3, 103-113.

[11] Lawler, E. L., "On Scheduling problems with Deferral Costs," *Management Science*, 11, 2 (1964), 280-288.

[12] Luss, H. and D. R. Smith,"Resource Allocation Among Competing Activities: A Lexicographic Minimax Approach," *Operations Research Letters*, 5, 5 (1986), 227-231.

[13] Nemhauser, G. L. and L. A. Wolsey, *Integer and Combinatorial Optimization*, John Wiley, (1988).

[14] Penlesky, R. J., W. L. Berry and U. Wemmerlov, "Open Order Due Date Maintenance in MRP Systems," *Management Science*, 35, 5 (1989), 571-584.

[15] Tang, C. S., "A Max-Min Allocation Problem : It's Solution and Applications," *Operations Research*, 36, 2 (1988), 359-367.